Report No. 5675

PACKET RADIO NETWORK FINAL REPORT

Karen Seo

Prepared by:

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238

Prepared for:

Dr. Barry Leiner and Dr. Robert Kahn
Information Processing Techniques Office
Defense Advanced Research Projects Agency

TABLE OF CONTENTS

1. INTRODUCTION

This report describes the work accomplished under contract MDA 903-80-C-0206, which covered the last three years of the Packet Radio program.

Over the past ten years, BBN has taken a leading role in the development of Packet Radio technology. In fulfilling this role we became the system architects of the Packet Radio program, creating the Tier and Multistation designs as well as several of their predecessors. At a more local level, we designed SPP, the reliable end-to-end protocol used within the network, and we have significantly contributed to the development of the TCP protocol used for Internet traffic. We have documented much of our work in the conference papers enumerated in chapter 2.

In addition, we have implemented the station software, which controls the routing and distributed name server function in the multistation environment. To insure adequate system integration and testing of our software, we have taken a leading role in the SRI testbed by specifying network configurations and tests to be run, by implementing network analysis programs, and by leading the system level debugging of the network. This last role has fallen naturally upon us because when something doesn't work, the problem is first noted at the top and must be traced downward.

Our acceptance of the role of system debugger has contributed to inter-contractor relations as we provided debugging tools, test scenarios, and on-call advice for any network implementation problem.

Much of the research performed under this contract was published as BBN reports and conference papers. The following chapter provides abstracts of the technical documents.

2. TECHNICAL DOCUMENTS

During the Packet Radio program, much of the research was documented first as Packet Radio Temporary Notes (PRTN) and later published as conference papers. An additional source of information on the station software is contained in the operator's manuals. This chapter contains abstracts for the technical documents that BBN has produced for this contract.

2.1 PRTNs

The following list of Packet Radio Temporary Notes documents the history of BBN's involvement in the development of the Packet Radio network and the gradual evolution of the current software.

PRTN 278, "New PRN Device ID Policy," presents a redefinition of the device identifier space which permits graceful exportation of the gateway function out of the station into a separate machine.

PRTN 279, IEN 120, "Internet Routing and the Network Partition Problem," describes an extension to routing design in the Internetwork environment, capable of routing to the correct partition of partitioned networks.

PRTN 280, "Transfer Points," gives details of the transfer

point, or "incremental," routing, which will permit very long routes and will support multiple station operation.

PRTN 281, "Multistation Design Specification," specifies how multiple station operation will work in the Packet Radio Network.

PRTN 285, "Configuring the PDP-11/34 Station for Auto-Restart," is a guide to reconfiguring station hardware and its related effects on operational procedure. It was prepared especially for use at Fort Bragg.

PRTN 288, "Congestion at the PR Hop Level," addresses likely inefficiencies in the transmission of packets from one Packet Radio unit to the next under the current protocol, and suggests particular changes to the hop protocol.

PRTN 289, "Multistation and CAP6 Routing Implementation Specification," incorporates comments received and negotiated agreements extending PRTN 281, resulting in a specification for implementers.

PRTN 289, rev 1, "Multistation and CAP6 Routing Implementation Specification," is intended to complete the details required for implementation of CAP6 and to provide some clarification of network algorithms. A number of technical

modifications were also made. This revision represents a joint effort with Rockwell International.

PRTN 289, rev 2, "Multistation and CAP6 Routing Implementation Specification," represents the completed CAP6.0 design and implementation details. This revision also has updated descriptions of the partially overlapping subnets, CAP6.1, but does not attempt a complete description.

PRTN 290, "Utilizing Internet Routes as Expressways Through Slow Nets," discusses routing traffic through additional, faster networks to reduce delay in delivering the traffic.

PRTN 292, IEN 146, "Flying Packet Radios and Network Partitions," applies the network partition solution, first introduced in PRTN 279, to the problem of airborne Packet Radio units which may be in contact with multiple ground networks simultaneously.

PRTN 295, "Delivering IPR Software to Station Disk," describes the procedure presently used to deliver IPR software to the Packet Radio station disk and documents the necessary support programs and operating environment.

PRTN 300, "CAP7: Station to Stationless Network Interface,"

proposes a major modification to stationless network design which is intended to provide better route assignments with fewer packet transmissions. This PRTN also explores a possible station to stationless network interface based upon the new stationless design.

PRTN 303, "Flooding in Multistation Routing," describes the flooding algorithm used in multistation route setup and explains some of the system requirements for successful implementation.

PRTN 304, "The Multistation Implementation," describes the modules used to perform the multistation function of route setup. The paper also discusses methods of handling errors in route selection and of speeding up route requests.

PRTN 310, "Degradable Performance in Packet Switching Networks," analyzes the behavior of an overloaded packet-switching node with respect to discarding and/or blocking traffic.

2.2 Manuals

This section contains a list of the manuals that can be referred to for information on how to operate the station software as well as for a detailed description of the routing

algorithms used.

"CAP 5.3 Labeler Operator's Guide" documents several new features of the station's Labeler process.

"CAP 5.3 Labeler" describes the operation and internal algorithms employed in this version of the station Labeler process, which receives network status information via the efficient Performance Data Packet mechanism and periodically improves assigned routes and labels.

"CAP 5.6 Labeler" describes the algorithms used to control the CAP5.6 network and to provide routing. Included is the revised Operator's Guide, which documents the new set of available commands

"CAP 6.0 Labeler" describes the algorithms used to control the CAP6.0 network and to provide routing. Included is the revised Operator's Guide, which documents the new set of operator commands.

"CAP 8 Station Notebook" contains the following information on the CAP8 station.

- Overview

- o Operator's guide for the station software
- o Information related to doing a system release
- o Descriptions of how each station process works
- o Manual on XNET, the cross-net debugger
- o Documentation on ELF, the station's operating system
- o Addenda on 11-BCPL, the language used in the station software
- o Trouble report form
- o Instructions for setting station identification numbers
- o Description of the station diagnostics

2.3 Conference Papers

The research that was done under this contract was presented at several conferences. This section contains abstracts of those papers.

William MacGregor, Jil Westcott, and Michael Beeler, "Multiple Control Stations in Packet Radio Networks," presented at MILCOM'82.

Abstract: This paper presents a design implemented by the Packet Radio Project for control of large networks. The network is built on a carrier-sense, multiple-access broadcast channel and is populated with mobile store-and-forward nodes known as Packet Radio units, or PRs.

Until recently, the Packet Radio network operated with centralized routing controlled by one node, the station. To increase the size of the network and to provide redundant control a multiple station design, "multistation," was necessary.

The PRs gather local connectivity information which is then forwarded to the controlling nodes (stations). Each station is responsible for up to 40 PRs located near it in connectivity and uses these PRs to communicate with other stations.

Questions important to the development of multistation include. How are Packet Radios matched with controlling stations? How do stations find and communicate with each other? How are long routes crossing the borders of many stations determined, and how is their successful setup insured? How is control traffic minimized?

Jil Westcott, BBN, and John Jubin, Rockwell International, "A Distributed Routing Design for a Broadcast Environment," presented at MILCOM'82.

<u>Abstract</u>: In this paper, we present the need for distributed network control in a Packet Radio environment, the criteria used for its evaluation, and an overview of the network algorithms chosen to maintain routing information and to route user traffic.

The routing algorithm, known as tiered rings, functions by building a distributed tree of shortest path routes to each Packet Radio in the network. The information kept by each node grows only linearly as the network expands. The design uses the radio channel efficiently by providing shortest path routing and by taking advantage of the information available in the broadcast environment.

Jil A. Westcott, "Issues in Distributed Routing for Mobile Packet Radio Networks," presented at COMPCON'82.

Abstract: This paper explores issues in distributed routing control for store-and-forward networks and emphasizes the needs of mobile Packet Radio networks. It includes a comparison of the proposed design for distributed routing in the Packet Radio network and the routing protocol used by the ARPANET, a wire network.

In particular, the paper describes proposed modifications to the distributed routing algorithm developed for the Packet Radio network. The primary change is in the handling of route failure due to the creation of route loops or a dead end arising from loss of connectivity. The purpose of the proposed change is to speed restoration of the route by making use of additional information in the neighborhood of the failure. This is

especially important in a mobile environment, because of the high frequency of altered connectivity.

Michael Beeler, "Degradable Performance in Packet Switching Networks," presented at COMPCON'82.

__Abstract__: A packet-switching node becomes overloaded when traffic arrives faster than the node can service it. We model such a node and examine its behavior. We identify two regions of overloaded operation, one where some traffic is discarded and one where additionally some traffic is blocked. We argue that discarding traffic is important, and identify parameters a designer may set to tailor the operation to his particular network application.

3. MEETINGS AND NEGOTIATIONS

3.1 Meetings

During this program, members of the BBN Packet Radio project participated in the following meetings.

- o Packet Radio Working Group (PRWG) meeting, SRI, December 4-6, 1979.
- o Internet meeting, SRI, February 4-6, 1980.
- o Multistation Implementers Technical Design meeting, BBN, March 6-7, 1980.
- o Second Multistation Implementers Technical Design meeting, BBN, April 30-May 1, 1980.
- o PRWG meeting, Fort Bragg, May 6-8, 1980.
- o Internet meeting, MIT, week of May 12, 1980.
- o CAP6 Design Review, SRI, January 26-27, 1981.
- o Internet meeting, ISI, January 28-30, 1981.
- o Packet Radio Working Group Meeting, SRI, April 14-16, 1981.
- o Internet Working Group Meeting, COMSAT, June 1-3, 1981.
- o Measurement Design Review, Washington, D.C., June 4-5, 1981.
- o Internet Technical Meeting, DARPA, July 8, 1981.
- o Packet Radio Working Group Meeting, SRI, September 1-3, 1981.
- o CAP7 Design Review Meeting, Rockwell, October 8-9, 1981.
- o Packet Radio Working Group Meeting, SRI, December 15-18, 1981

- Internet Meeting, BBN, March 24-25, 1982.
- Packet Radio Planning Meeting, Lake Tahoe, Nevada, April 19-21, 1982.
- Packet Radio Implementation Meeting, DARPA Headquarters, June 8-9, 1982.
- Packet Radio Working Group Meeting, San Diego, California, June 22-24, 1982.

3.2 Negotiations

Also during this period, we led negotiations with the other Packet Radio contractors to achieve or approach closure on the following design and implementation issues.

- Design and implementation specifications of periodic point-to-point route improvement.
- Packet Radio checksumming algorithm design, including our running of execution speed tests in a PDP-11 station.
- Assessment of using seven-hop routes at Fort Bragg, resulting in the conclusion reported to ARPA that seven-hop routes be adopted.
- Study of the impact of upgrading transfer point design for support of multistation; impact was deemed small, and packet formats were specified which include all fields necessary for both transfer points and multistation.
- Assessment of secondary storage capacity (presently disk) required by the station, as related to the future use of bubble memory; resolved that bubbles are reasonable, especially because of reduction in expected size, weight, power, and possibly hardware failure rate.
- Discussion of the congestion and load problems experienced at Fort Bragg. We recommended that they install the CAP5.6 software to improve performance.

- o Proposal of a new route loop detection method designed to overcome loops undetectable because of action by the duplicate filter table.

- o Design and documentation of the CAP6.0 network which resulted in a joint revision of PRTN 289, "Multistation and CAP6 Routing Implementation Specification," with Rockwell International.

- o Investigation of aberrant PR down load characteristics. Study of packet transcripts revealed simultaneous loading by two stations. BBN maintained a correspondence with Rockwell over several iterations of the subsequent design changes and agreed with the final proposal.

- o Discussion of the stationless design (CAP7) to conclude the final details.

4. SOFTWARE DEVELOPMENT

4.1 Station

This section describes the work done on the different versions of the station software and some of the other development work that was done to build the current Packet Radio network.

4.1.1 CAP5 development

We completed the CAP5 software, which handles a single station network with distributed store-and-forward algorithms and point-to-point route assignment by the station.

- o Implementation of the CAP5.2 Labeler was completed. This Labeler periodically requests source/destination pairs from PRs and, if routes are present, transmits the new best routes.
- o Release to SRI of a Labeler which periodically updates labels in addition to routes.
- o Addition of command to print the text of label packets.
- o Addition of command to print the text of new or repaired point-to-point routes.
- o Addition of command to print all routes improved during periodic updating, including source and destination identifiers.
- o Addition of command to print SPP connection state changes, especially helpful in tracking a connection-handling bug.
- o Modification of Labeler to use seven-hop routes. Although by far the largest changes were to the Labeler,

other station processes (connection, XRAY debugger, and PR down line loader) were also affected. This change, constituting CAP5.3, was tested and delivered to SRI, Collins, and Fort Bragg networks.

- o Completion of the mini-station programming and testing, and delivery of the CAP5 mini-station software to SRI in time for the GALLANT EAGLE military exercise.

4.1.2 CAP6 development

We completed and delivered to SRI the CAP6.2 software for multistation operation. This was an upgrade to a multistation design which allowed for longer routes and larger networks.

- o Implementation, debugging, and testing of the completely overlapping multistation protocol.
- o Implementation of the partially overlapping design.

4.1.3 CAP7 development

In August 1981, BBN presented PRTN 300, "CAP7: Station to Stationless Network Interface," proposing both a new stationless network design and the outline of a design which interfaces station and stationless areas. At the CAP7 Design Review meeting in October, the group decided to adopt the "tiered ring" design proposed by BBN

The CAP7 design is described in the paper, "A Distributed Routing Design for a Broadcast Environment," by Westcott and Jubin, which was presented at MILCOM'82, at Boston, Mass., in

October 1982.

4.1.4 CAP8 development

We completed and implemented the CAP8 design, which provides further integration of the combination of multistation and stationless protocols developed in CAP6 and CAP7 software.

4.1.5 IPR down load software

Over the years the IPR download software underwent several changes.

- The PRLOAD process, and also the system to dump PR load images onto station disk, were extensively changed to support down line loading of both Experimental (EPR) and Improved (IPR) versions of the Packet Radio unit.
- There was a successful demonstration of down line loading PR operating system code, CAP protocol, and assorted diagnostics into IPRs.
- Later the new PRLOAD system was transferred to Collins and SRI.
- We also added a command to print the status of down loading and a command to list the names of PR load images present on the station disk.

4.1.6 Assimilation of PDP-11 TCP4

This software was received from SRI and integrated into the station.

4.1.7 Bootstrap

We implemented a bootstrap using 96-bit leaders for use on the station and minigateway when attached to a Port Expander. This was also required for attachment to any ARPANET port after January 1981. It was delivered to SRI.

4.1.8 Conversion to SPP2 protocol

We converted all the latest versions of station software from the SPP protocol to the SPP2 protocol, and coordinated this conversion and testing of the results with Collins. Our part of this task was principally a rewrite of the station's connection process, but the Labeler, XRAY, and PRLOAD also needed significant modification. There were several benefits:

- Performance Data Packets (PDPs) arriving at the station in a burst are all accepted, buffered, and serviced. Previously, the delay to open a new "listening" connection after receipt of almost every PDP caused others arriving on the heels of the first to be discarded. Use of SPP2 thus provides more complete control information and prevents channel burdening from retransmitted PDPs.

- The 20-second timeout, necessary to prevent SPP oscillation, no longer exists. Thus connections to PRs can be reestablished as quickly as necessary after a transaction is completed. This achieves smoother, more responsive, more efficient network control. Assignment of new, improved, or repaired routes will never be held up by the 20-second timeout, as occurred in some cases in the past.

- Switching to the simpler SPP2 allowed simplification of programs, resulting in a savings of station memory space needs. The connection process shrank in size by about

50 percent; some reduction in other processes also occurred.

- o The simpler handling and the space savings permit the support of more simultaneous connections between the station and PRs (or other devices). Only ten words of storage must now be allocated per connection to a PR, so on the order of a hundred connections can be maintained.

- o The conversion to the simpler SPP2 made implementation of multistation operation easier.

4.1.9 LSI-11 Mini-Station

We converted the Packet Radio station software to run on the LSI-11-based Mini-Station, which is now designed to use one RXV-21 floppy disk and one Bubble-Tec bubble storage unit.

- o We wrote drivers for the 1822 interface (manufactured by ACC) and for the RXV-21 floppy disks. In addition, BBN completed the changes to ELF, bootstraps, PRLOAD, disk drivers, and the SPP2 process, which proved necessary to use the floppy disks.

- o We assisted SRI in the installation and debugging of the bubble storage units.

4.2 TCP and IP

Work was done in several areas for the Transmission Control Program (TCP) and Internet Protocol (IP).

- o Changes to host TCP/IP implementation to support "long" 1822 leaders and IMP numbers exceeding 6 bits.

- o Modification of the TOPS-20 internet user queue facility to support single-port protocols such as the new IP4 version of the XNET debugger.

- o Integration of monitor-based Internet Protocol layer version 4 (IP4), Transmission Control Program version 4 (TCP4), TCP Virtual Terminals (TVTs), and TELNET into the operating systems TENEX, TOPS-20 release 4 as supplied by DEC, and the BBN version of TOPS-20 release 4. This included testing, debugging and distribution of the software.

- o Changes to host IP software to support multiple network interfaces (Multinet). This included installation of algorithms to read a gateway descriptor file as the system comes up, determine whether gateways are up or down, and permit operator notification when the gateway descriptor file is found to be out of date.

- o Installation, testing, and debugging of the ICMP protocol. This included improvement of the routines to listen to host "redirect" and "network dead" messages.

- o Implementation and testing of a File Transfer Process (FTP) for TOPS-20 and TENEX sites based on version 4 of TCP and IP.

- o Problem investigation and debugging assistance, eg., installation of extensive packet tracing tools, for groups including ISI, MIT, and BBN:

  . checksum errors in reset packets

  . flow control problems such as "silly window syndrome" and retransmission strategies

  . connection closing problems in user TELNET

  . connection refusals

  . free storage management problems

  . lock management problems

  . TCP-to-user signals problems

  . TELNET virtual terminals (TVTs) problems

4.3 Gateway

Another area in which BBN did significant work was gateway software.

- o Delivery of minigateway software to execute in a true minigateway, an LSI-11. Previous version had run in PDP-11 machines because of lack of LSI-11 availability.
- o Revised the gateway software to include fragmentation and option processing.
- o Concluded gateway development and transferred the technology to another BBN group (under separate contract to DARPA) for standardization, operational deployment, and long-term maintenance.

5. SUPPORT SERVICES

5.1 Station software

There were several areas in which BBN supplied support for the station software.

- o We were responsible for the maintenance and delivery of CAP5 and CAP6 versions of station software to testbeds at Fort Bragg, SRI, Rockwell and SAC.

- o Through cooperation with personnel at Collins Radio, the effort to bring their PDP-11 up as a station was completed. This effort included consultation and writing diagnostics on a disk specially prepared for Collins. Having a station at Collins significantly speeded software checkout of PR code and permitted joint debugging sessions where we connected to their station via the XNET debugger. This facility was instrumental in fast turnaround of PR software fixes and the implementation of seven-hop routes (CAP5.3), SPP2 protocol, IPR down line loading, IPR performance improvements, and IPR memory (PROM) retrofit (these last two being Collins' tasks).

5.2 Packet Radio software

We provided reformatting and delivery services for Collins' PR code at all sites.

5.3 Gateway software

We maintained and delivered a considerable volume of software. For example, the gateways that fell under our responsibility included:

- o four ELF-based gateways on Packet Radio networks
- o three ELF-based gateways on the SATNET
- o two mini-gateways on Packet Radio networks
- o four SATNET mini-gateways.

In addition, we had responsibility for delivery of gateway software to:

- o University College London
- o Fort Bragg
- o SRI.

5.4 Demonstrations

BBN supported both field activities at the testbeds and special demonstrations.

- o National Telecommunications Conference (NTC) This effort included modifying the XNET debugger to load the gateway at COMSAT, modifying the gateways to support changes to the catenet configuration, and preparation of a stable internet environment on the test system, BBNF, for use as a demonstration host site.
- o HELBAT This involved LSI-11 Mini-Station development, which allowed mini-stations to be deployed in the HELBAT (Human Engineering Laboratory Battalion Automation Test) exercise at Fort Sill, Oklahoma, in October 1981.
- o Gallant Eagle We provided software support and advice for the deployment of Packet Radio mini-stations in the GALLANT EAGLE exercise, in the week of March 1, 1982.
- o TCP We provided TOPS-20 TCP/IP assistance for the above exercises.

5.5 Debugging Packet Radio network

In general BBN lead the system level debugging of the network. for example. we participated in the hunt for the cause of bad PR identifiers (IDs). This was a problem that had been plaguing the SRI network since spring of 1979. Our contribution included consultation and two special modifications to the Labeler process. One printed new IDs. which were often bad ones. and the PR forwarding them as they arrived at the station. The second permitted the operator to control point-to-point routes by optionally disabling the fixing of bad routes

5.6 TCP and IP user assistance and support

Along with developing and distributing the new TCP/IP software for the TOPS-20 and TENEX, BBN provided the following general user-level support.

- o Assisted users at BBN. COMSAT. and ISI in the testing of TELNET Programs for the HP3000, TAC. and RT-11
- o Assisted users in the installation of IP and TCP in both TOPS-20 and TENEX systems
- o Revised and updated the documentation describing the user-level interface to IP and TCP
- o Investigated problems and provided debugging assistance to groups at ISI, MIT. BBN, and other user sites.